

On Breaching Enterprise Data Privacy Through Adversarial Information Fusion

Srivatsava Ranjit Ganta, Raj Acharya

Dept of Computer Science, Penn State University
University Park, PA, USA
ranjit@cse.psu.edu
acharya@cse.psu.edu

Abstract—Data privacy is one of the key challenges faced by enterprises today. *Anonymization* techniques address this problem by *sanitizing* sensitive data such that individual privacy is preserved while allowing enterprises to maintain and share sensitive data. However, existing work on this problem make inherent assumptions about the data that are impractical in day-to-day enterprise data management scenarios. Further, application of existing anonymization schemes on enterprise data could lead to adversarial attacks in which an intruder could use *information fusion* techniques to inflict a privacy breach. In this paper, we shed light on the shortcomings of current anonymization schemes in the context of enterprise data. We define and experimentally demonstrate *Web-based Information-Fusion Attack* on anonymized enterprise data. We formulate the problem of *Fusion Resilient Enterprise Data Anonymization* and propose a prototype solution to address this problem.

I. INTRODUCTION

Data privacy is one of the key challenges faced by enterprises today. Sensitive individual-specific information such as customer data, employee data etc are maintained and used for various purposes. Several instances of data privacy breaches [1] in the recent past have resulted in financial as well as reputation losses for enterprises. *Anonymization* techniques address this problem by *sanitizing* sensitive data such that individual privacy is preserved while allowing enterprises to maintain and share sensitive data. Recently, there has been a lot of work [2] [3] [4] [5] [6] on data anonymization schemes. These techniques can be broadly classified into two types:

- **Partitioning based anonymization schemes** : The first class of techniques guarantee privacy by partitioning the data such that an adversary cannot uniquely identify the individuals falling in each partition. The basic ideology behind these techniques is *blending in the crowd* which guarantees that an individual or entity cannot be distinguished from a minimum number of other people. *K*-anonymity [2], *l*-diversity [4] and other work in this line [7] achieve partitioning through *generalization* and *suppression* techniques. On the other hand, techniques such as [8], [9] achieve this by *clustering* the data. Partitioning based solutions are mainly applied to *non-interactive* scenarios where the data needs to be *published/released* after anonymization.
- **Perturbation based anonymization schemes** : The other

class of techniques guarantee data privacy by adding *noise* to the sensitive data and thus preventing identification. Solutions in this category can be further classified based on whether the setting considered is interactive or not. Solutions such as [5] [6] add noise to perform specific data mining tasks in a non-interactive setting. More recent solutions such as [10] add randomized noise in an *interactive* setting where-in the particular function to be evaluated on the data is known *a priori*.

In this paper, we consider a *non-interactive* setting where the data needs to be *released/published*. We focus on partitioning based schemes as they are readily applicable to generic databases including data with categorical attributes. Table I depicts a typical individual-specific data considered in partitioning based anonymization literature. Observe that there exists a *classification* of data attributes (as shown in Table I) into three different types:

- 1) **Identifier Attributes**: Attributes carrying explicit identifiers such as *Name*, *SSN* etc.
- 2) **Quasi Identifier Attributes**: Attributes that could indirectly lead to identification of individuals in the database such as *Age*, *Zipcode* and *Gender* etc. These are also sometimes referred to as *Non-Sensitive* attributes.
- 3) **Sensitive Attributes**: Attributes carrying the sensitive information about the individuals such as *Disease*, *Income* etc.

Based on this classification, existing solutions assume that the *Identifier Attributes* in the database are stripped *prior* to the anonymization process. This was under the implicit assumption that the identifier attributes were necessary neither for the release nor for the intended purpose of the release. We believe that this assumption is too restrictive and is even impossible in some scenarios where the presence of explicit identifiers is *necessary* for the intended purpose of the anonymized release [11]. Consider the following scenario:

Enterprise Data - Example : Table II depicts a customer database in a typical financial institution. The data contains names of all the customers along with certain *non-sensitive* and *sensitive* information. The non-sensitive attributes are: *Investment Volume Index (Invst Vol)* to indicate the volume of investment (number of shares traded etc.) made by the

Identifiers		Quasi Identifiers			Sensitive
Name	SSN	Zipcode	Age	Nationality	Condition
Alice	111-111-1111	13053	28	Russian	AIDS
Bob	222-222-2222	13068	29	American	Flu
Christine	333-333-3333	13068	21	Japanese	Cancer
Robert	444-444-4444	13053	23	American	Meningitis

TABLE I
SENSITIVE DATABASE

customer in the past, *Investment Amount Index (Invst Amt)* to indicate the amount of investment (amount involved in previous trades etc.) made by the customer in the past, *Customer Valuation (Valuation)* to indicate the assigned value of the customer. The only *sensitive* attribute, *Customer Personal Income (Income)*, corresponds to the customer's personal income. Databases such as this are an integral part of enterprises and are maintained and used for key operations everyday. In this paper, we shall refer to them as *Enterprise Databases*.

The internal *release* of such data along with explicit identifiers (*Customer Names*) is a necessity for several enterprise operations such as accounting, record keeping etc. However, at the same time, such a release should not compromise the privacy of sensitive information (*Customer Personal Income*). Note that trivial solutions such as removal of identifiers or use of pseudonyms are not viable in such scenarios. The key properties here are:

- The inclusion of identifying information is necessary for the *release* to serve the intended purpose.
- Sensitive data disclosure should not be compromised even in the presence of explicit identifiers.

Name	Invst Vol	Invst Amt	Valuation	Income
Alice	8	7	4	91,250
Bob	5	4	4	74,340
Christine	4	5	5	75,123
Robert	9	8	9	98,230

TABLE II
ENTERPRISE DATA

In the enterprise database scenario described above, anonymizing data using existing techniques falls short in providing adequate protection against adversarial attacks. This is because existing techniques [2] [3] [4] make an assumption that *Identifier Attributes* are stripped *prior* to the anonymization process. Consider the possibility in which an adversary (possibly an *insider*) is given (or otherwise acquires) access to the anonymized release of an enterprise database. Now, the adversary can use the identifiers present in the release to collect *auxiliary* information about the individuals present in the database from a multitude of sources such as the web (homepages, blogs etc). The adversary could then *fuse* the auxiliary information with the anonymized release to estimate sensitive data.

Web-Based Information-Fusion Attack : Consider the enterprise data example described earlier as shown in Table II.

One way to internally release this table is to remove the customer salary information and publish the non-sensitive data as it is. The problem with this approach is that one can estimate the sensitive data based on the non-sensitive information present in the release. The solution is to anonymize the non-sensitive information and remove the sensitive information. Table III shows the anonymized release of this data using partitioning based anonymization scheme such as *K*-anonymity proposed by Sweeney et al. [2]. We use *K*-anonymization as a representative of partitioning based solutions for data anonymization as other solutions in this category produce similar results.

Name	Invst Vol	Invst Amt	Valuation	Income
Alice	[5-10]	[5-10]	[1-5]	-
Bob	[5-10]	[1-5]	[1-5]	-
Christine	[1-5]	[1-5]	[1-5]	-
Robert	[5-10]	[5-10]	[5-10]	-

TABLE III
ANONYMIZED ENTERPRISE DATA

Table III is now deemed *safe* and is released internally within the enterprise. Now, consider the scenario in which an adversary employee *Bob* is granted access to this anonymized release. Note that the release does not give *Bob* the sensitive information i.e customer personal income data. However, he has access to non-sensitive information such as the customer valuation, investment volume etc. *Bob's* goal is to use the anonymized release to estimate the customer personal income values. To achieve this, he uses the customer names present in the release to search for additional information about the customers available on the web which will help him estimate their personal income. For example, he collects information about the customer's *Employment*, *Property Holdings* etc. Example of such data collected from the web is shown in Table IV. Now, by *fusing* this information with the anonymized release the adversary can estimate the sensitive customer personal income information. In this example, let's say the income range for

Name	Employment	Property Holdings
Alice	CEO, Deutsche Bank	3560
Bob	Manager, Verizon	1200
Christine	Assistant, NYU	720
Robert	CEO, Microsoft	5430

TABLE IV
AUXILIARY DATA COLLECTED BY THE ADVERSARY

all the customers is [\$40000 - \$100000] and could be divided into three classes *Low* [\$40000 - \$60000], *Medium* [\$60000 - \$80000], and *High* [\$80000 - \$100000]. Now, consider the customer *Robert*. With an estimated valuation falling in the highest range [5-10], *Bob* concludes that *Robert* falls into the highest income category [\$80000 - \$100000]. By looking at his employment and property holdings (and possibly other auxiliary information), *Bob* can further improve his estimate and conclude that *Robert* falls into upper category [\$90000 - \$100000] of the *High* income class. Based on this, he estimates that *Robert's* salary is the average of range [\$90000 - \$100000] i.e \$95000. This example demonstrates, how, by using the auxiliary information obtained from the web an adversary could obtain a close estimate of *Robert's* actual income. Although in the above example the attacker uses his understanding of the data to *fuse* the anonymized release with web data, in reality, he could use various *Information Fusion* techniques for this purpose. Information Fusion is a well-studied paradigm in which multiple data sources are used to improve knowledge extraction.

In the attack demonstrated above, an adversary with access to anonymized enterprise data gleans auxiliary information from the web and uses information fusion techniques to inflict a privacy breach. In this paper, we refer to such an attack as *Web-Based Information-Fusion Attack* on enterprise data. This is illustrated in Figure 1. Note that this attack is an example of an attack-model in which a *human-in-the-loop* inflicts a privacy breach.

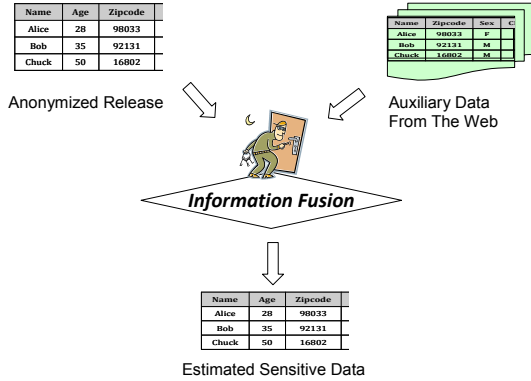


Fig. 1. Web-Based Information-Fusion Attack

A. Contributions and Organization

In this paper, we demonstrate the shortcomings of existing anonymization schemes when applied to enterprise data through the *Web-Based Information-Fusion Attack*. Our main contribution is the formulation of *Fusion Resilient Enterprise Data Anonymization* problem. We propose an iterative scheme to find an optimal anonymization that offers maximum protection against such attacks for a given dataset. The rest of the document is organized as follows: Section 2 provides the related work to this problem. Section 3 elaborates on the Web-Based Information-Fusion Attack and discusses the assumptions made regarding the attack. In Section 4, we formulate the

problem of Fusion Resilient Enterprise Data Anonymization. We then present our solution strategy to address the problem through incremental anonymization in Section 5. Section 6 presents experimental results by demonstrating the attack on a real data set and presenting the prototype solution. Section 7 provides the conclusion and future work.

II. RELATED WORK

Data privacy has received a lot of attention from both computer science and statistical research communities. In statistical literature, studies on data confidentiality [12] [13] propose the use of matrix masks for anonymizing data. In the computer science literature, several recent studies [2] [3] have been done in the context of *K*-anonymity. Ferrer [9] proposed heuristic algorithms for optimal *K*-anonymization on quantitative data. Several problems with *k*-anonymity based partitioning techniques have been studied in [4] [7] and others. In [4], Machanavajjhala et al. pointed to the possibility of attacks on *k*-anonymized data because of lack of diversity in the sensitive values corresponding to each partition. Later, in [7], Li et al. provided an argument that *l*-diversity is neither a sufficient nor a necessary condition to guard against attacks on *k*-anonymized data. They proposed a scheme in which the distribution of sensitive values within each partition should not be far from the distribution of sensitive values in the original data.

One of the primary challenges in data anonymization is to take into consideration the *auxiliary* information (also called *external knowledge*, *background knowledge* or *side information*) that an adversary can glean from other channels. Recent work on partitioning based techniques [4] [14] [15] has attempted to define adversary's *background knowledge* and possible privacy breach based on this. Martin et al [14] provide a first formal treatment of adversarial background knowledge. They propose a language for expressing the adversary's knowledge based on conjunctive propositions. More recently, Chen et al. [15] have attempted to fill this gap, by proposing an extension to the same language based framework. However, these models do not consider auxiliary information obtained *using identifying information present in the anonymized release*.

On the other hand, there has been some work [16] [17] [18] on addressing the problem of anonymizing sequential releases. The problem here is to ensure that the current release of a particular data set does not lead to a disclosure with respect to previous releases on the same data set. Orthogonal to these works, in [19] Wong et al prove that adversary's knowledge of the anonymization algorithm could lead to a privacy breach. In [20], Aggarwal et al. pose the problem of adversarial rule mining attack on anonymized data. Our work is critically different from these studies as we consider inferential attribute disclosure based on *Information Fusion* using external information sources.

III. WEB-BASED INFORMATION-FUSION ATTACK

A. Information Fusion

In this paper, we use *fuzzy inferencing* to build an Information Fusion system. This section provides a brief introduction to fuzzy inferencing and how it can be used by the adversary to fuse the anonymized release with web-based auxiliary information.

Fuzzy Inference is a well-studied paradigm based on *fuzzy logic*, *fuzzy if-then rules* and *fuzzy reasoning*. Basically, it provides a mechanism to *map* a set of *inputs* to a set of *outputs* using a set of *rules*. We refer the reader to [21] for an introduction to fuzzy inference systems. The first step involved in creating a fuzzy inference system is to determine the *inputs* and *outputs*. In the web-based information-fusion attack, the inputs include all the data attributes available to the adversary through: 1. The anonymized release and 2. The auxiliary data collected through the web. In our running example from Section 1, the attributes *Investment Volume Index*, *Investment Amount Index*, *Customer Valuation* from the anonymized release in Table III form the first half of inputs to the information fusion system. The attributes *Employment*, *Property Holdings* collected from the web form the second half of inputs. The output consists of single attribute, *Customer Personal Income*, which the adversary intends to estimate. In the second step, the adversary defines *fuzzy-set definitions* for each of the input and output attributes. He then uses domain knowledge to formulate a set of *rules* mapping the input fuzzy sets to the output fuzzy sets. Figure 2 illustrates the system.

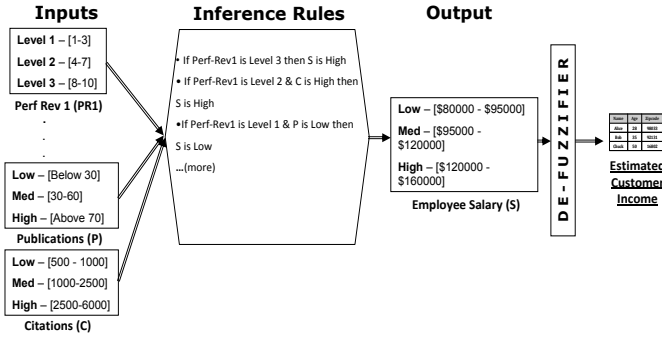


Fig. 2. Fuzzy Inference System

B. Attacker Capability

We assume that the intruder is an *insider* who is given or otherwise acquires access to the anonymized data. Thus, the intruder has access to individual identifiers that can be used to index into the web and other data sources. The intruder is assumed to have the domain knowledge about the data to perform information fusion.

IV. PROBLEM FORMULATION

In this section we formulate the problem of *Fusion Resilient Enterprise Data Anonymization* to address web-based information-fusion attacks. Since it is not possible to quantify the amount of auxiliary information the adversary can collect,

it is not practical to completely prevent such attacks. However, by estimating the auxiliary information that an adversary could collect, we can *minimize the extent of privacy breach* in case of such an attack. This forms the primary goal of our problem formulation: For a given sensitive dataset, we need to find an anonymization such that the release causes minimum breach in case of a fusion attack. On the other hand, one of the important factors involved in data anonymization is the *utility* of the release [22] [3]. The utility of an anonymized release is a measure of usefulness of the release for the intended purpose such as a specific task to be performed on the data Ex. Classification etc. Several standard measures such as [22] have been proposed in the literature to compute data utility. Hence, the secondary goal of our problem formulation is to maximize data utility. With these goals in hand, we proceed to formulate the overall goal as follows:

Let $P = \{p_{ij}\}_{m \times n}$ be a sensitive private dataset defined over a finite set of attributes $\{P_1, P_2, \dots, P_n\}$.

Let $Q = \{q_{ij}\}_{r \times s}$ be the auxiliary data gathered by the intruder from the web over a set of attributes $\{Q_1, Q_2, \dots, Q_s\}$. Now, let P' be a *candidate* anonymization of P .

Let F be an *information fusion* system that takes in P' and Q as inputs and produces \hat{P} , an estimate of P .

Let U be a measure of utility of P' .

Goal : The goal of *Fusion Resilient Enterprise Data Anonymization* is to compute a P' from P such that:

- 1) P' is *resilient* to *Web-based Information Fusion Attacks*.
- 2) The utility U offered by P' meets the release requirements.

To formulate the problem based on the above goal, we need to quantify the *resilience* to web-based information-fusion attacks. We define this using the following definitions:

Definition 1 Dissimilarity ($D_1 \circ D_2$) For two datasets D_1 and D_2 representing the same set of individuals and the same set of attributes, $D_1 \circ D_2$ is a measure of dissimilarity between them.

For two datasets $\{D_1\}_{m \times n}$ and $\{D_2\}_{m \times n}$ representing the same set of individuals, we compute the dissimilarity using mean square distance D_1 and D_2 :

$$D_1 \circ D_2 = \frac{1}{m} * Tr((D_1 - D_2)^T (D_1 - D_2))$$

where m is the total number of records in each database and $Tr(A)$ of a matrix A is the trace of A , i.e the sum of the elements of the main diagonal.

As defined earlier, \hat{P} is an estimate of P made by the adversary based on a candidate release P' and web-based auxiliary data Q using the information fusion system F .

$$\hat{P} = F(P', Q)$$

In order for privacy of P to be protected, the dissimilarity between P and the estimate made by the adversary, \hat{P} , needs to be *large*. The more the dissimilarity $P \circ \hat{P}$, the better protected P is. Also, the dissimilarity between P and \hat{P} *quantifies* the *protection* offered by the corresponding P' against information

fusion attacks. Based on this, we now define a *Fusion Resilient Anonymization* as:

Definition 2 Fusion Resilient Anonymization An anonymization P' of a given sensitive data P is resilient to fusion attacks if the dissimilarity $(P \circ \hat{P})$ between \hat{P} and P is above a certain threshold value T_p .

So, for a candidate anonymization P' to be a *safe* release, the corresponding $(P \circ \hat{P})$ needs to be above a certain threshold value T_p . It is obvious to note that, among all the possible anonymizations (P' 's) that satisfy this property, the one that has maximum value of $(P \circ \hat{P})$ offers maximum protection. So, for the anonymization P' to offer maximum resilience to web-based information fusion attacks, the *dissimilarity* $(P \circ \hat{P})$ needs to be maximized.

Recall that in addition to maximizing the protection against information-fusion attacks, the utility of the release (U), should be maximized. Let W_1 and W_2 be the weights assigned by the publisher for privacy protection against information fusion attacks and data utility respectively. Now, the final objective can be stated as a *weighted sum of protection and utility* of the form:

$$W_1 * (P \circ \hat{P}) + W_2 * U$$

Now, the problem can be stated as,

Problem : Given a private dataset P , web-based data Q and an information-fusion system F , find the fusion resilient anonymization P' that maximizes $H = W_1 * (P \circ \hat{P}) + W_2 * U$, where \hat{P} represents the estimate of P based on P' and Q using F .

In order to solve the above *optimization problem*, we need to find the *optimal* anonymization P' in the *solution space* containing all possible anonymizations P' 's that satisfy the fusion-resilient-anonymization property defined earlier. One way to look at this *solution space* is to consider the set of all anonymizations possible by anonymizing P to different *levels*. Note that the definition of *Anonymization Level* depends on the specific anonymization scheme to be employed. For example, in K -anonymization, the value of k represents the anonymization level. The more the value of k is, the more the anonymization level. As mentioned in Section 1, in our work, we use K -anonymization as the basic anonymization scheme. For a given dataset P , let i denote the anonymization level and P'_i denote the release obtained by anonymizing P to level i . We use the *discernibility metric* defined in [22] to measure the utility of a k -anonymized data set. The metric can be mathematically stated as follows.

$$C_{DM}(g, k) = \sum_{\forall |E| \geq k} |E|^2 + \sum_{\forall |E| < k} |D| * |E|$$

where E refers to the *clusters* or *equivalence classes* of the data set induced by k -anonymization of g using the value k . The reader is referred to the original paper for further details. Based on the above definition, let the utility of P'_i be denoted by U_i . The optimization function H can now be defined based on anonymization level i as:

$$H_i = W_1 * (P \circ \hat{P}_i) + W_2 * U_i$$

Let T_u be the minimum utility required for the release. Now, the above generic problem statement can be instantiated as:

Problem Statement: Find $P'_{i_{opt}}$, such that

$$H_{i_{opt}} = \max_{\forall i} H_i$$

where, $(P \circ \hat{P}_i) \geq T_p$ and $U_i \geq T_u$.

V. SOLUTION

In this section we propose a simple iterative algorithm to find the *fusion resilient anonymization* for a given sensitive dataset. The strategy is to take any basic anonymization scheme such as k -anonymization and *incrementally* anonymize the data. The level of anonymization is increased in steps (increase k in steps), until the utility of the release falls below a threshold. In each step, the web-based fusion attack is simulated to find whether the resulting candidate anonymization offers enough protection. If yes, the candidate anonymization is retained, otherwise it is discarded. This results in a set of all *candidate* anonymizations present in the *solution space*. We then search for the optimal anonymization level that offers the maximum weighted sum of protection and utility. Figure 3 illustrates our approach.

Algorithm 1 presents this solution in procedural format as FRED_Anonymization (Fusion Resilient Enterprise Data Anonymization). The algorithm uses the Basic_Anonymization procedure that takes a sensitive data and level of anonymization as inputs and produces an anonymization of the input data to the corresponding anonymization level. For this, any basic anonymization algorithm such as the ones proposed in [9] [3] can be used to generate a k -anonymization. Note that in case of k -anonymization the minimal level of anonymization is achieved by using the value $k = 2$. The algorithm uses the Basic_Anonymization procedure to anonymize the sensitive data for increasing values of the anonymization level (*level*). The stopping condition for this loop is achieved when the utility of anonymized result (P') denoted by U_{level} falls below the threshold T_u . In each iteration, the algorithm simulates an information fusion attack to produce the estimate an adversary could obtain (\hat{P}_{level}). The *dissimilarity* between the estimated values \hat{P}_{level} and the original values P is computed using the procedure Dissimilarity_Measure which takes two datasets as input and outputs the dissimilarity value as described in Section 4. At this point, the dissimilarity is compared against a threshold value T_p to check if the anonymization offers enough protection against information fusion attacks. If yes, the weighted sum of dissimilarity and utility is computed and stored as $H(i)$. Finally, the algorithm searches for the anonymization level i that offers the maximum value for the weighted sum of protection and utility H_{max} . The anonymization $P'_{i_{opt}}$ corresponding to H_{max} is the fusion resilient anonymization of the original data that offers maximum weighted protection as well as utility.

VI. EXPERIMENTAL RESULTS

In this section we present experimental results by demonstrating the *web-based information-fusion attack* on a real-life

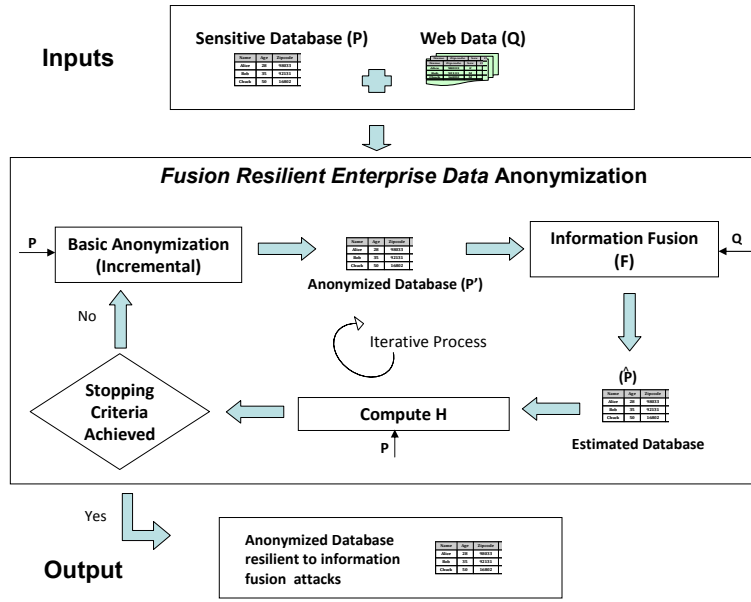


Fig. 3. Fusion Resilient Enterprise Data Anonymization

Algorithm 1 FRED_Anonymization

```

1:  $P \leftarrow$  Sensitive Data
2:  $Q \leftarrow$  Web Data
3:  $F \leftarrow$  Information Fusion System
4:  $T_p \leftarrow$  Protection Threshold
5:  $T_u \leftarrow$  Utility Threshold
6:  $W_1 \leftarrow$  Protection Weight
7:  $W_2 \leftarrow$  Utility Weight
8:  $level \leftarrow -1$ 
9:  $i \leftarrow 0$ 
10: repeat
11:    $level \leftarrow level + 1$ 
12:    $P'_{level} \leftarrow \text{Basic\_Anonymization}(P, level)$ 
13:    $\hat{P}_{level} \leftarrow F(P'_{level}, Q)$ 
14:    $P \circ \hat{P}_{level} \leftarrow \text{Dissimilarity\_Measure}(P, \hat{P}_{level})$ 
15:    $U_{level} \leftarrow \text{Utility}(P'_{level})$ 
16:   if  $(P \circ \hat{P}_i) \geq T_p$  then
17:      $H(i) \leftarrow W_1 * (P \circ \hat{P}_{level}) + W_2 * U_{level}$ 
18:      $i \leftarrow i + 1$ 
19:   end if
20: until  $U_{level} \geq T_u$ 
21:  $i_{max} \leftarrow i - 1$ 
22:  $H_{max} \leftarrow H(0)$ 
23: for  $i = 1$  to  $i = i_{max}$  do
24:   if  $H(i) \geq H_{max}$  then
25:      $i_{opt} \leftarrow i$ 
26:   end if
27: end for
28: return  $P'_{i_{opt}}$ 

```

dataset. The goals here are to quantify the information gained by the adversary through information fusion and demonstrate the FRED_Anonymization algorithm.

A. Setup

The sensitive data (P) is collected from a real-life enterprise (a public university) and contains salary information and performance review numbers of the employees (faculty). The employee *Salary* is the *sensitive* attribute while the performance review numbers are the *non-sensitive attributes*. The data is anonymized (P') so as to *suppress* all of the salary information and k -anonymize the non-sensitive attributes using *microaggregation* based k -anonymization proposed in [9]. The external data (Q) is collected from the employee web pages and external links from there. Based on domain knowledge, we formulate a simplistic set of knowledge rules to fuse P' and Q and build a fuzzy inference system to estimate the employee salary as illustrated in Figure 2. All the rules are assigned uniform weights.

All the experiments were implemented using Matlab on a PC with Intel Pentium 4 (1.8GHz) processor and 1GB of RAM running Microsoft Windows XP.

B. Information Gain

Our first study aims to quantify the *information gain* obtained by the attacker in estimating the sensitive data P , by introducing web-based auxiliary information Q . Consider the adversary's knowledge of the original data at two stages 1. *Before* information fusion, and 2. *After* information fusion. Recall that to start with, the adversary has access to the anonymized release P' . The adversary then collects Q and fuses this with P' to obtain \hat{P} . So, before performing information fusion, the adversary's (best) knowledge about the original data is the anonymized version itself, i.e P' (in the absence of Q). In

this case, we have the dissimilarity between the original and the adversary's estimate $(P \circ \hat{P}) = (P \circ P')$. Figure 4 plots this $(P \circ P')$ for increasing values of k . It is not surprising to observe that the *dissimilarity* increases as k increases, since the *level* of anonymization increases with k . After performing information fusion, the adversary obtains \hat{P} by fusing P' with Q using F . Figure 5 plots this $(P \circ \hat{P})$ for increasing values of k . Notice that $(P \circ \hat{P})$ is lesser than $(P \circ P')$ for all values of k . In other words, the estimate made by the attacker (\hat{P}) after information fusion is closer to (P) than when compared to the estimate available before information fusion (P'). The difference between $(P \circ P')$ and $(P \circ \hat{P})$ is precisely the amount of information gained by the adversary through information fusion. Hence, the *Information Gain* G of the adversary is the difference between the closeness of the estimates available before and after information fusion.

$$G = (P \circ P') - (P \circ \hat{P})$$

Figure 6 plots G for increasing values of k . It is interesting to observe that G does not necessarily increase with k . This implies that as the level of anonymization increases, the information gained by the attacker decreases. The reason for this is that as the level of anonymization increases, the input (P') to the information fusion system gets worse and thus forces the system to output incrementally bad estimates.




Fig. 4. Before Information Fusion $(P \circ P')$

C. Optimal Anonymization

We now study the fusion resilient enterprise data anonymization that leads to maximum weighted sum of protection and utility as formulated in Section 4. We use the *discernibility metric* defined in [22] to measure the utility of a k -anonymized data set. The basic idea here is to assign each data sample (or vector) a *cost* based on the number of data vectors it is indistinguishable from, or in other words, the size of the cluster it falls into. If the cluster size it falls into is greater than k , then the cost assigned is equal to the size of the cluster. If the cluster size is less than k , then the cost is much severe (since it does not adhere to the definition of




Fig. 5. After Information Fusion $(P \circ \hat{P})$




Fig. 6. Information Gain (G)

k -anonymity) and is equal to the product of the size of the whole data set and the size of the cluster.

$$C_i = \begin{cases} |E|^2 & \text{if } |E| \geq k \\ |D| * |E| & \text{otherwise} \end{cases}$$

Using this definition, we define the utility of the data set $U = \{u_{i1}\}_{m \times 1}$ as a column matrix where each entry is the inverse of the cost assigned to the corresponding data point.

$$u_{i1} = 1/C_i$$

To show how utility of the release varies with increasing level of anonymization (increasing values of k), we calculate the utility of the entire release using the discernibility definition [22] as:

$$C_{DM}(k) = \sum_{\forall |E| \geq k} |E|^2 + \sum_{\forall |E| < k} |D| * |E|$$

$$U_k = 1/C_{DM}(k)$$

Figure 7 plots U_k for increasing values of k . It is straightforward to observe that utility of data decreases as k increases. The goal now is to find the optimal k value such that the resulting anonymization offers maximum weighted sum of

privacy protection and utility formulated as:

$$H = \frac{1}{m} * Tr((P \circ \hat{P})^T W_1 (P \circ \hat{P})) + \frac{1}{m} * Tr(U^T W_2 U)$$

We establish the threshold values for protection and utility as $T_p = 3.075$ $T_u = 0.0018$ based on experimental observations. For these threshold values, we obtain the solution space of $k = 7$ to 14. We assign equal weights to privacy protection and utility i.e $W_1 = W_2 = 0.5$, $W_i = 0.5_{m \times n}$ i.e . Based on this setup, Figure 8 plots H for increasing values of k within the solution space. By running an optimization for the maximum value of H , we obtain the result $k = 12$. This is the optimal k value that provides the maximum weighted sum of protection and utility.




Fig. 7. Utility U_k




Fig. 8. Weighted Sum Of Protection And Utility H_k

VII. CONCLUSION

In this paper, we establish two problems encountered in privacy preserving enterprise data management:

- 1) Enterprise Data anonymization involves minimizing data disclosure in the presence of explicit individual-identifier information.
- 2) Existing anonymization techniques fall short in protecting enterprise data privacy in case of adversarial information fusion.

We defined the *Web-Based Information-Fusion Attack* wherein an adversary uses information fusion techniques to fuse anonymized data with publicly available information from the web to inflict a privacy breach. Our experimental demonstration of the attack present the practicality and easiness with which such attacks might lead to revelation of sensitive data. We formulate the problem of finding a *fusion resilient data anonymization* and propose a simple solution to address this problem. While it is not possible to entirely prevent fusion based privacy attacks, one can minimize the extent of breach possible through intelligent data anonymization.

REFERENCES

- [1] R. Agrawal, J. Kiernan, R. Srikant, and Y. Xu, "Hippocratic databases," in *28th Int'l Conf. on Very Large Databases (VLDB)*, Hong Kong, 2002. [Online]. Available: citeseer.ist.psu.edu/agrawal02hippocratic.html
- [2] P. Samarati and L. Sweeney, "Protecting privacy when disclosing information: k-anonymity and its enforcement through generalization and suppression," *Technical Report, CMU*, 1998.
- [3] K. LeFevre, D. J. DeWitt, and R. Ramakrishnan, "Mondrian multidimensional k-anonymity," in *ICDE*, 2006, p. 25.
- [4] A. Machanavajjhala, J. Gehrke, and D. Kifer, "l-diversity: Privacy beyond k-anonymity," in *ICDE*, 2006.
- [5] R. Agrawal and R. Srikant, "Privacy-preserving data mining," in *Proc. of the ACM SIGMOD Conference on Management of Data*. ACM Press, May 2000, pp. 439–450. [Online]. Available: citeseer.ist.psu.edu/agrawal00privacypreserving.html
- [6] A. Evfimievski, R. Srikant, R. Agrawal, and J. Gehrke, "Privacy preserving mining of association rules," in *Proc. of 8th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining (KDD)*, July 2002. [Online]. Available: citeseer.ist.psu.edu/evfimievski02privacy.html
- [7] N. Li, T. Li, and S. Venkatasubramanian, "t-closeness: Privacy beyond k-anonymity and l-diversity," in *ICDE*, 2007, pp. 106–115.
- [8] G. Aggarwal, T. Feder, K. Kenthapadi, S. Khuller, R. Panigrahy, D. Thomas, and A. Zhu, "Achieving anonymity via clustering," in *PODS*, 2006, pp. 153–162.
- [9] J. Domingo-Ferrer, "Practical data oriented microaggregation for statistical disclosure control," *IEEE TKDE*, 2002.
- [10] C. Dwork, "Differential privacy," in *ICALP (2)*, 2006, pp. 1–12.
- [11] S. R. Ganta and R. Acharya, "Adaptive data anonymization against information fusion based privacy attacks on enterprise data," in *SAC*, 2008.
- [12] G.T.Duncan and Lambert, "Disclosure limited data dissemination," *Journal of American Statistical Association*, 1986.
- [13] G.T.Duncan, "Inferential disclosure -limited microdata dissemination," *Proceedings of the Survey Research Methods Section, American Statistical Association*, 1990.
- [14] D. J. Martin, D. Kifer, A. Machanavajjhala, J. Gehrke, and J. Y. Halpern, "Worst-case background knowledge for privacy-preserving data publishing," in *ICDE*, 2007, pp. 126–135.
- [15] B.-C. Chen, R. Ramakrishnan, and K. LeFevre, "Privacy skyline: Privacy with multidimensional adversarial knowledge," in *VLDB*, 2007, pp. 770–781.
- [16] K. Wang and B. C. M. Fung, "Anonymizing sequential releases," in *KDD*, 2006, pp. 414–423.
- [17] J.-W. Byun, Y. Sohn, E. Bertino, and N. Li, "Secure anonymization for incremental datasets," in *Secure Data Management*, 2006, pp. 48–63.
- [18] X. Xiao and Y. Tao, "M-invariance: towards privacy preserving republication of dynamic datasets," in *SIGMOD Conference*, 2007, pp. 689–700.
- [19] R. C.-W. Wong, A. W.-C. Fu, K. Wang, and J. Pei, "Minimality attack in privacy preserving data publishing," in *VLDB*, 2007, pp. 543–554.
- [20] C. Aggarwal, J. Pei, and B. Zhang, "On privacy preservation against adversarial data mining," in *SIGKDD*, 2006.
- [21] B. Kosko, *Neural Networks and Fuzzy Systems*. Prentice Hall.
- [22] R. Bayardo and R. Aggarwal, "Data privacy through optimal k-anonymization," in *ICDE*, 2005.